



Task-Level Authoring for Remote Robot Teleoperation

Emmanuel Senft^{1*}, Michael Hagenow², Kevin Welsh¹, Robert Radwin³, Michael Zinn², Michael Gleicher¹ and Bilge Mutlu¹

¹Department of Computer Sciences, University of Wisconsin–Madison, Madison, WI, United States, ²Department of Mechanical Engineering, University of Wisconsin–Madison, Madison, WI, United States, ³Department of Industrial System and Engineering, University of Wisconsin–Madison, Madison, WI, United States

OPEN ACCESS

Edited by:

Stefanos Nikolaidis,
University of Southern California,
United States

Reviewed by:

Xuning Yang,
Carnegie Mellon University,
United States
Yi-Shiuan Tung,
University of Colorado Boulder,
United States

*Correspondence:

Emmanuel Senft
esenft@wisc.edu

Specialty section:

This article was submitted to
Human-Robot Interaction,
a section of the journal
Frontiers in Robotics and AI

Received: 09 May 2021

Accepted: 03 September 2021

Published: 27 September 2021

Citation:

Senft E, Hagenow M, Welsh K,
Radwin R, Zinn M, Gleicher M and
Mutlu B (2021) Task-Level Authoring
for Remote Robot Teleoperation.
Front. Robot. AI 8:707149.
doi: 10.3389/frobt.2021.707149

Remote teleoperation of robots can broaden the reach of domain specialists across a wide range of industries such as home maintenance, health care, light manufacturing, and construction. However, current direct control methods are impractical, and existing tools for programming robot remotely have focused on users with significant robotic experience. Extending robot remote programming to end users, i.e., users who are experts in a domain but novices in robotics, requires tools that balance the rich features necessary for complex teleoperation tasks with ease of use. The primary challenge to usability is that novice users are unable to specify complete and robust task plans to allow a robot to perform duties autonomously, particularly in highly variable environments. Our solution is to allow operators to specify shorter sequences of high-level commands, which we call task-level authoring, to create periods of variable robot autonomy. This approach allows inexperienced users to create robot behaviors in uncertain environments by interleaving exploration, specification of behaviors, and execution as separate steps. End users are able to break down the specification of tasks and adapt to the current needs of the interaction and environments, combining the reactivity of direct control to asynchronous operation. In this paper, we describe a prototype system contextualized in light manufacturing and its empirical validation in a user study where 18 participants with some programming experience were able to perform a variety of complex telemanipulation tasks with little training. Our results show that our approach allowed users to create flexible periods of autonomy and solve rich manipulation tasks. Furthermore, participants significantly preferred our system over comparative more direct interfaces, demonstrating the potential of our approach for enabling end users to effectively perform remote robot programming.

Keywords: human-robot interaction, end-user programing, teleoperation, robotics, remote robot control, user study

1 INTRODUCTION

Effective teleoperation of robots—broadly, a remote human controlling a robot at a distance (Niemeyer et al., 2016)—is critical in scenarios where automation is impractical or undesirable. When a person operates a remote robot, they must acquire sufficient awareness of the robot's environment through sensors and displays, be able to make decisions about what the robot should do, provide directions (control) to the robot, and evaluate the outcomes of these operations. These challenges have been addressed with a wide range of interfaces that span a continuum of levels of

autonomy (Beer et al., 2014), ranging from direct control where the operator drives the moment-to-moment details of a robot's movements, to asynchronous control, where operators send complex programs to the robot to execute autonomously, e.g., space exploration where robots receive programs for a day's worth of activities (Maxwell et al., 2005). The choices of level of control provide different trade-offs to address the goals of a specific scenario. In particular, longer-horizon control offers better robustness to communication issues and provides long periods of idle time for the operator while the robot is executing the commands. However, it also limits the opportunities for the human to react to unexpected situations during the program execution and requires significant human expertise to design robust behaviors and advanced sensing skills for the robot. On the other hand, more direct control allows operators to react quickly and easily to uncertainty, but demands constant attention from the operator, often relies on dedicated hardware, and requires a fast and stable connection to ensure that the tight real-time loop between the operator and the robot is maintained.

Our goal is to provide effective telemanipulation for end-user applications, such as home care, light manufacturing, or construction. In such scenarios, high level robot autonomy of autonomy would be desirable, as this would reduce the operator's workload, however there remains situations where a fully autonomous behavior cannot be created. Users have domain knowledge, they can analyze the robot environment and determine appropriate actions for the robot, but they have no expertise in creating robot programs. Building a system supporting teleoperation for these novice users presents a number of challenges, the system needs to 1) be easy to use, 2) support active perception (Bajcsy, 1988), 3) support specification of robot behaviors adapted to the current state of the environment, and 4) allow for periods of autonomy. As we will detail in **Section 2**, current interfaces for teleoperation are often specifically tailored to highly trained operators or adopt a low level of autonomy. The former are not suited to novice users and the latter forces users to continuously provide inputs to the robot, reducing both the usability over extended periods of time and increasing the sensitivity to communication issues.

Our key idea is to use task-level authoring to enable the operator to control the robot by specifying semantically connected sequences of high-level (task-level) steps. This paradigm supports various lengths of program depending on available environment information, ranging from single actions to longer plans. For example, a robot might need to open a drawer with a specific label and empty it, however the robot does not have character recognition. The operator could use the robot to locate the appropriate drawer and then create a plan for the robot to open this specific drawer, remove all items in it, and then close it. Task-level authoring aims to offer more flexibility for the operator, allowing them both to specify long periods of autonomy when possible, but also have a more direct control when necessary to allow the operator to obtain the environmental awareness necessary to make longer plans.

We propose four principles to support effective telemanipulation by novices:

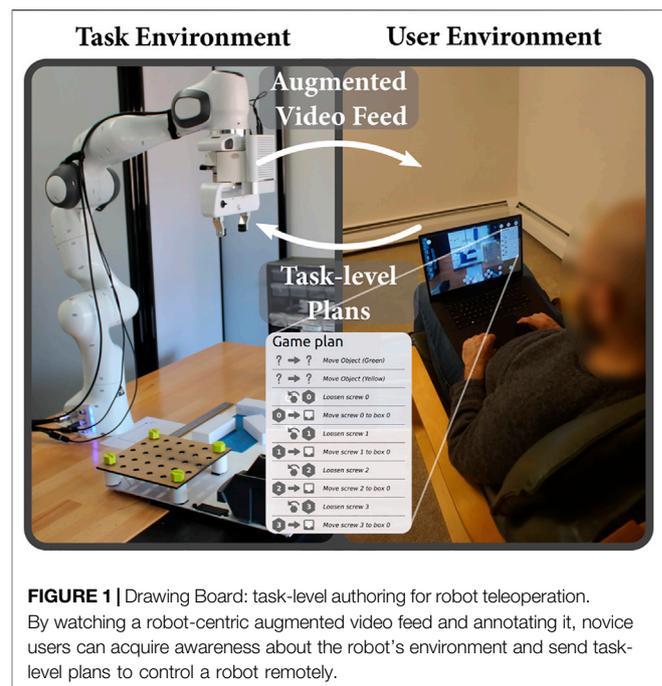


FIGURE 1 | Drawing Board: task-level authoring for robot teleoperation.

By watching a robot-centric augmented video feed and annotating it, novice users can acquire awareness about the robot's environment and send task-level plans to control a robot remotely.

- 1) Interleaving observation and planning: the stepwise nature of manipulation tasks allows phases of observing the environment to gain awareness with phases of acting on that information. Execution occurs asynchronously, allowing it to be robust against communication problems and providing idle time to the user. Users can assess the state of the environment, devise a short plan for the robot, execute it, and restart the process with the new state of the environment.
- 2) Controlling the robot at the action level: instead of controlling the robot motions, operators can select actions for the robot (e.g., pick-up, pull, or loosen). Such higher level of control allows participants to focus on the task that needs to be solved instead of the robot's kinematics or workspace geometry.
- 3) Providing a unified augmented reality interface: task specification can be accomplished from a viewpoint chosen by the user to be convenient, as part of their awareness gathering process. This process allows us to use a screen overlay-based augmented reality interface that aggregates all the required information for decision making on a single view also used to specify actions. This single view makes the programming easier for the users as all the important information are available in a single place.
- 4) Specifying actions graphically: the augmented-reality interface allows for details of operations to be specified and verified graphically in context, simplifying the interface further. Additionally, such graphical specification allows easily to generalize a plan to a group of objects of the same type.

We have prototyped these ideas in a system called Drawing Board after an artist's portable drawing board (**Figure 1**) and evaluated it in a user study with 18 participants. Our central contribution is to show that a task-level authoring approach can be applied to teleoperation to create a system that affords both ease-

of-use and asynchronous operation. In our study, remote operators (students with limited programming knowledge) were able to perform complex tasks, gaining the benefits of asynchronous operation (robustness to delays and opportunities for longer periods of idle time) with the ease-of-use and reactivity of more direct interfaces. Participants—some literally on the other side of the world—were able to teloperate the robot with little training and preferred our system compared to interfaces not embodying our principles. These findings show how the core choice of task-level authoring is supported by specific interface and implementation designs, yielding a system that meets our goals, allowing end users to remotely create short period of autonomy for robots.

2 RELATED WORK

Our work brings elements from the field of authoring and end-user programming to teleoperation.

2.1 Teleoperation

Fundamentally teleoperation refers to human control of robot actions, typically done remotely (i.e., the human and the robot are not collocated and the human can only perceive the robot's environment through artificial sensors and displays) (Niemeyer et al., 2016). With teleoperation, the question of the appropriate level of autonomy is important, especially in the presence of delay and partial situational awareness (Yanco et al., 2015; Niemeyer et al., 2016). Levels of autonomy form a continuum between direct control and long-term programs:

- 1) Direct control: low-level control where the operator is manually controlling all actions of the robot in real-time (e.g., remote surgery (Marescaux et al., 2001), military (Yamauchi, 2004)).
- 2) Semi-autonomy: the human operator intermittently controls robot actions where required and can parameterize higher-level actions that are executed autonomously by the robot (e.g., search-and-rescue, DARPA Robotics Challenge (Johnson et al., 2015)).
- 3) Teleprogramming: operators create programs defining actions and reactions to changes in the environment for the robot to execute over longer period of time (e.g., Mars rovers programmed every day for a full day of autonomy (Norris et al., 2005)).

Direct control has seen widespread use in the aerospace, nuclear, military, and medical domains (Niemeyer et al., 2016) as it allows operators to quickly react to new information. However, this type of teleoperation requires constant inputs from the operator and is highly sensitive to communications problems. Researchers have explored various methods to address this communication challenge. One direction of research involves optimizing the communication channel itself to reduce delay and allow the operator to have quick feedback on their actions (Preusche et al., 2006). Another method, shared control, seeks to make the process more robust to human error through means such as virtual fixture methods, which support the operator in

their direct manipulation task (Rosenberg, 1993) or alternating phases of teleoperation and autonomous operation (Bohren et al., 2013). Finally, a third alternative uses a virtual model of the workspace to provide rapid feedback to the user from simulation while sending commands to the robot (Funda et al., 1992).

On the other end of the spectrum, traditional programming for autonomy and teleprogramming provides only limited feedback to the operator about the robot behavior. Operators need to have complete knowledge about the task including all required contingencies, to create dedicated programs for each task. These programs must be robust enough to run autonomously for hours without feedback. Furthermore, the robot needs to have the sensing capabilities to capture and analyze every relevant information in the environment. This highly autonomous control method is especially useful where there are large time delays between the robot and the operator which prevents the operator from intervening in real-time, such as when controlling a rover on Mars (Maxwell et al., 2005).

A semi-autonomous robot is a middle ground between these two extremes: it can execute short actions autonomously, but relies on the human operator to determine a plan of action and provide the correct parameters for these actions. The human (or team of humans) can use the robot to actively collect information about the environment, and provides near real-time inputs to the robot. The DARPA robotic challenge explores this space. In this case, the robot can run parameterized subroutines while multi-person teams of highly trained operators analyze data from the robot and control it at various abstraction-levels (from joint angle to locomotion goal), including situations with unstable communication channels (Johnson et al., 2015). These subroutines can be parameterized by selecting or moving virtual markers displaying the grasping pose (Kent et al., 2020), robot joint position (Nakaoka et al., 2014), or using affordance templates (Hart et al., 2014). In a retrospective analysis, Yanco et al. (2015) highlight the training required for operating the robots during these trials, and reports that researchers should explore new interaction methods that could be used by first responders without extensive training.

One approach to simplify both awareness acquisition and control (two keys aspects in teleoperation) is to use monitor-based augmented reality—overlaying digital markers on views from the real world (Azuma, 1997). For example, Schmaus et al. (2019) present a system where an astronaut in the space station controlled a robot on earth using this technology. Their point-and-click interface presents the video feed from the head camera of a humanoid robot with outlines of the detected objects and menus around the video. When clicking on one of these objects, the system filters the actions that can be done on this object to propose only a small subset of possible actions to the operator. Similarly, Chen et al. (2011) propose a multi-touch interface when actions are assigned to gestures on a video feed displayed on a touch screen. This type of point-and-click or gesture interface allows the remote operator to gain awareness about the environment and simply select high-level actions for the robot to perform. Simulations can also be overlaid with markers that users can manipulate to specify the desired position of a robot or its end-effector (Hashimoto et al., 2011; Hart et al., 2014; Nakaoka

et al., 2014). However, despite these advances, little work has been done to explore and evaluate interfaces that allow naive operators to actively acquire awareness about the environment and create longer plans consisting of multiple actions.

2.2 Authoring

In the context of robotics, the term authoring refers to methods allowing end users to create defined robot behaviors (Datta et al., 2012; Guerin et al., 2015; Weintrop et al., 2018). The general process starts with a design period where an initial behavior is created, then the robot can be deployed in the real world and its behavior tested and refined in additional programming steps if needed. When the desired requirements are achieved, the authoring process is finished and the robot is ready to be deployed to interact autonomously. Authoring differs from classic programming in its focus on end users with limited background in computer sciences and seeks to address questions of how can these users design, or author, behaviors using modalities such as tangible interactions (Sefidgar et al., 2017; Huang and Cakmak, 2017), natural language (Walker et al., 2019), augmented- or mixed-reality (Cao et al., 2019a; Peng et al., 2018; Akan et al., 2011; Gao and Huang, 2019), visual programming environments (Glas et al., 2016; Paxton et al., 2017), or a mixture of modalities (Huang and Cakmak, 2017; Porfirio et al., 2019). Steinmetz et al. (2018) describe task-level programming as parameterizing and sequencing predefined skills composed of primitives to solve a task at hand. Their approach combines this task-level programming and programming by demonstration (Billard et al., 2008) to create manipulation behaviors.

While promising, classic authoring methods suffer from two limitations when applied to remote robot control. First, the authoring process is often considered as a single design step creating a fully autonomous behavior (Perzlyo et al., 2016; Cao et al., 2019b). This monolithic approach differs from teleoperation which assumes that human capabilities (sensory or cognitive) are available at runtime to help the robot successfully complete a task. Second, many authoring methods such as PATI (Gao and Huang, 2019) or COSTAR (Paxton et al., 2017) use modalities only available in situations where the human operator and the robot are collocated (e.g., kinesthetic teaching, tangible interfaces, or *in-situ* mixed reality). For example, teach pendants—which are interfaces provided by manufacturers of industrial robots—are designed to be used next to the robot and often require the operator to manually move the robot. Consequently, while available to end users, such methods are not possible to use remotely. Our work is in the line of Akan et al. (2011), who used augmented reality to specify plans for a robotic arm. However, our premise is that to enable novice users to teleoperate robots, active perception (i.e., environment exploration) and behavior specification should be interleaved and coupled through a single simple interface, and that manipulation of graphic handles is a powerful way to specify parameters for actions.

3 DESIGN

To allow non-expert users to control robots remotely, we propose a system rooted in task-level authoring which allows users to navigate the live environment and specify appropriate robot

behaviors. The following sections and **Figure 2** detail the key concepts of the system topology.

3.1 Interleaving Observation and Planning

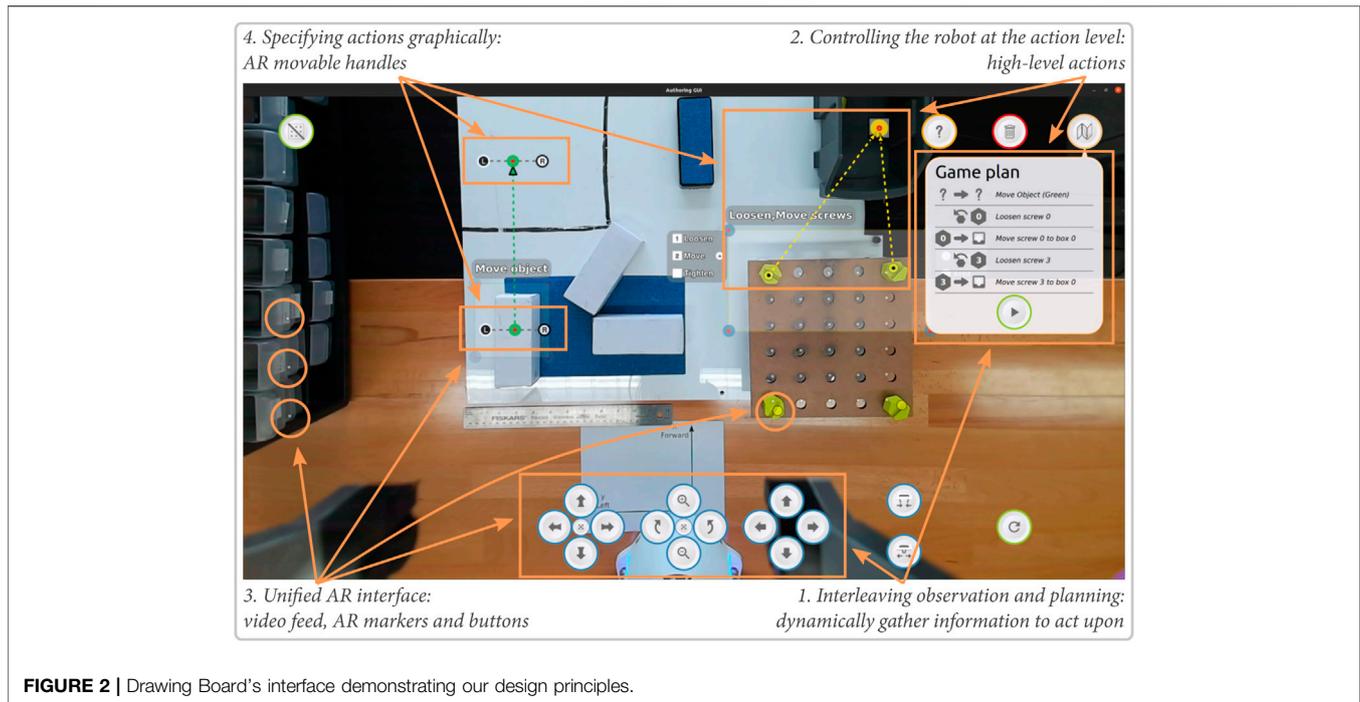
Specifying full execution plans for a robot would allow to reduce the operator workload during plan execution, but requires significant expertise in robotics and highly capable robots. To allow end users to create adaptable periods of autonomy for the robot, we propose to use a task-level authoring approach. This approach simplifies the programming process by allowing the programmer to break tasks into sequences of high level actions based on what they observe at the moment. Users can chain together actions to create flexible periods of autonomy, adapted to their knowledge of the situation. For example, a set of actions may consist of grabbing a set of bolts in an area and moving them into a set grid pattern to fasten a structure. If the user is unsure what action is required next or if something unexpected occurs, the task-level authoring approach allows the user to explore the environment and create new programs based on the outcome of previous actions and new information.

Controlling robots at the task level creates a number of opportunities for end-user teleoperation; it allows the human to remain in the decision loop to provide necessary expertise, while maintaining an asynchronous workflow. Such design allows end-users to alternate between observing the environment, specifying robot actions, and executing sequences of commands. Operators can specify short actions to explore the environment by moving the robot camera, acquiring awareness, and selecting an appropriate view point to author task plans. Then, once they have gathered enough information about the environment to know their next actions, they can schedule a longer plan consisting on multiple actions to solve the current part of the task. This process can be repeated as much as needed which allows for plans to be tailored to the current state of the environment. The asynchronous execution also provides robustness to communication instability. The inclusion of the operator in the control loop takes away the complexity of teleprogramming by having the human making complex perceptions and decisions. Thus, it keeps the benefits of direct control without the requirement of a tight and stable control loop and maintain the benefits of asynchronous control without requiring to create complex programs and plan ahead for unknown future.

3.2 Controlling the Robot at the Action Level

As mentioned in **Section 2.1**, teleoperation levels of control covers a spectrum from direct control to teleprogramming. Direct control can afford ease of use when the user is provided with intuitive input device (Rakita et al., 2018), however it requires minute control from the operator and is very sensitive to delay.

As shown in Schmaus et al. (2019), controlling a robot at the action level provides a number of advantages for teleoperation. First, as actions are executed using a local control loop, it allows to be robust to delays in communication. Second, it is intuitive for users, new operators can pick-up the system easily without requiring the user to possess any knowledge about robotics and control. Nevertheless, controlling solely at the action level



suffers from some limitations. When using a single actions, even when users know what the robot should do over the next few actions, they have to specify an action, wait for it to be executed, specify the next action, and repeat, which can be suboptimal for the user. Additionally, similarly to any high-level control scheme, any action not in the robot vocabulary cannot be executed.

3.3 Providing a Unified, Augmented Reality Interface

Similar to some other robot authoring interfaces (Schmaus et al., 2019; Walker et al., 2019), our approach uses augmented reality (AR) to simplify perception and action specification. The interface is composed of a unified monitor-based AR interface showing a live camera view of the robot's environment augmented with digital markers representing detected objects (see **Figure 2**). The camera is mounted directly on the robot's end-effector for viewpoint flexibility and registration. The interface is overlaid with a canvas where the operator can design robot behaviors. This paradigm is consistent with research which shows that the most intuitive way to communicate information to an untrained operator is through vision (Yanco et al., 2004). More complex information such as the detected object pose and the environment point cloud are used to parameterize robot behavior behind the scenes, but hidden from the user's display.

3.4 Specifying Actions Graphically

One challenge in designing an interface for novice end users is to simplify the specification of complex manipulations. In programming, classic ways to set parameters are through sliders and numbers. Numerical parameter-setting allows

greater precision, but can be unintuitive for users. Instead, our interface design leverages graphical representations whenever possible and minimizes required user input.

Our interface uses visual and interactive representations, mapped onto the augmented video feed, that enable users to parameterize predefined actions by manipulating these graphical representations. For example, to move a known object to a known positions, the interface creates anchors that can be moved by the user. Then, the interface will display an arrow from the starting point in the video to the goal point, visually representing the action in context. The interface uses 2D affordances throughout, as this is consistent with the 2D representation in video. 6D locations are inferred from the 2D interface based on environment information. Additionally, through graphical localization, a series of actions on a specific object can be generalized to nearby objects of the same category.

Our interface design only exposes high-level actions to the user (e.g., move, tighten, pull). The local robot controller decomposes these high-level actions into series of lower-level actions and translates them into primitives to reach the desired robot behavior. The user only has to specify the minimum fields required to execute the task and graphical specification allows to specify multiple parameters at the same time and in an intuitive manner.

4 IMPLEMENTATION

4.1 System

Following the considerations detailed in the previous section, we implemented Drawing Board, a prototype focused on enabling

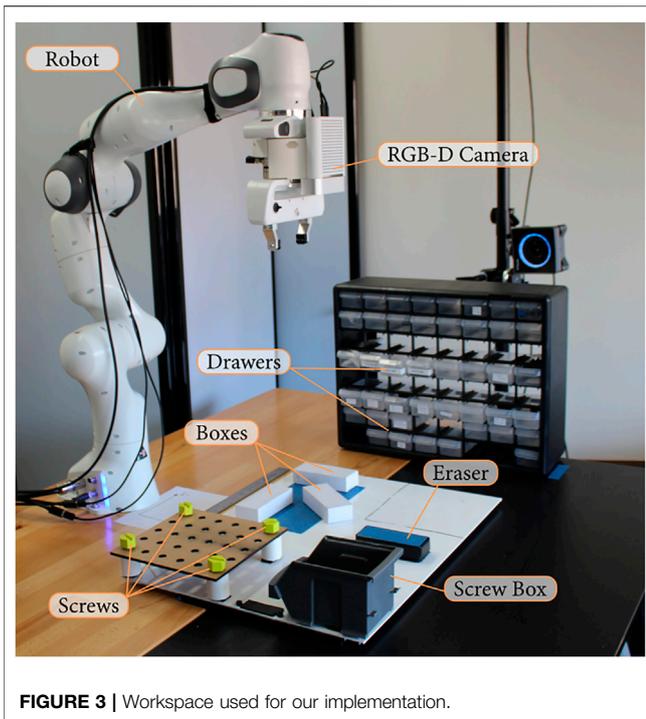


FIGURE 3 | Workspace used for our implementation.

users with little programming experience to operate a robot remotely. The interface was designed to be served from a traditional laptop or desktop display and focuses on controlling a single robotic manipulator. Our implementation integrates a collaborative robot (Franka Emika Panda) outfitted with an ATI Axia80-M20 6-axis force torque sensor and a Microsoft Azure Kinect providing both a 2D image and 3D point cloud. The camera is placed at the end-effector to allow for the greatest flexibility in camera position. The components of the system communicate using ROS (Quigley et al., 2009) with logic nodes implemented in Python, the graphical user interface in QML, and the low-level control in C++.¹

For facilitating precise interaction with the environment, we implemented a hybrid controller to have more control over the forces applied by the robot when doing precise manipulation such as pulling a drawer. The hybrid control law follows an admittance architecture where interaction forces are measured from the force-torque sensor and resulting velocities are commanded in joint space via pseudo-inverse based inverse kinematics.

We also leverage the Microsoft Azure Kinect depth sensor to observe the environment. Objects are first localized in the scene by feeding the color image to Detectron2 (Wu et al., 2019), which provides a high fidelity binary pixel mask for each detected object. Once the object is localized, a GPU-accelerated Hough transform is used to register the known triangle mesh with each instance. This pipeline allows us to achieve 6D object pose estimation, which can then be used to provide the user with semantically

correct actions as well as inform the robot motion plan. Our system also uses a number of predefined points of interest that represent the position of known static objects in the workspace. These known points are used as reference for the robot and to filter the position of objects detected by the live pose estimation pipeline.

4.2 Workspace

We applied our system to the workspace shown in **Figure 3**. This workspace guided our implementation, but the system can be adapted to other tasks or interactive objects. This workspace is composed of a number of drawers on the left of the robot with known positions. The middle of the workspace contains three white boxes above a blue area and a blue eraser, which are not detected by the robot object recognition system. The right part of the workspace contains a grid with holes and a screw box with known positions, and the grid can contain screws which are detected by the vision system.

The current prototype includes the following actions:

- Pulling and pushing the drawers;
- Picking, placing, and moving detected objects (e.g., screws) and undetected objects (e.g., boxes);
- Tightening and loosening the screws;
- Wiping an area.

By having general actions such as pull or move our system can adapt easily to other objects or different locations and orientations for these objects.

4.3 Interface

The default interface layout shows the video feed augmented with markers showing the detected or known points of interest (see **Figure 2**). The camera view is cropped to fill the full screen while showing clearly the robot's finger to allow users to know the gripper's status (open, closed, full).

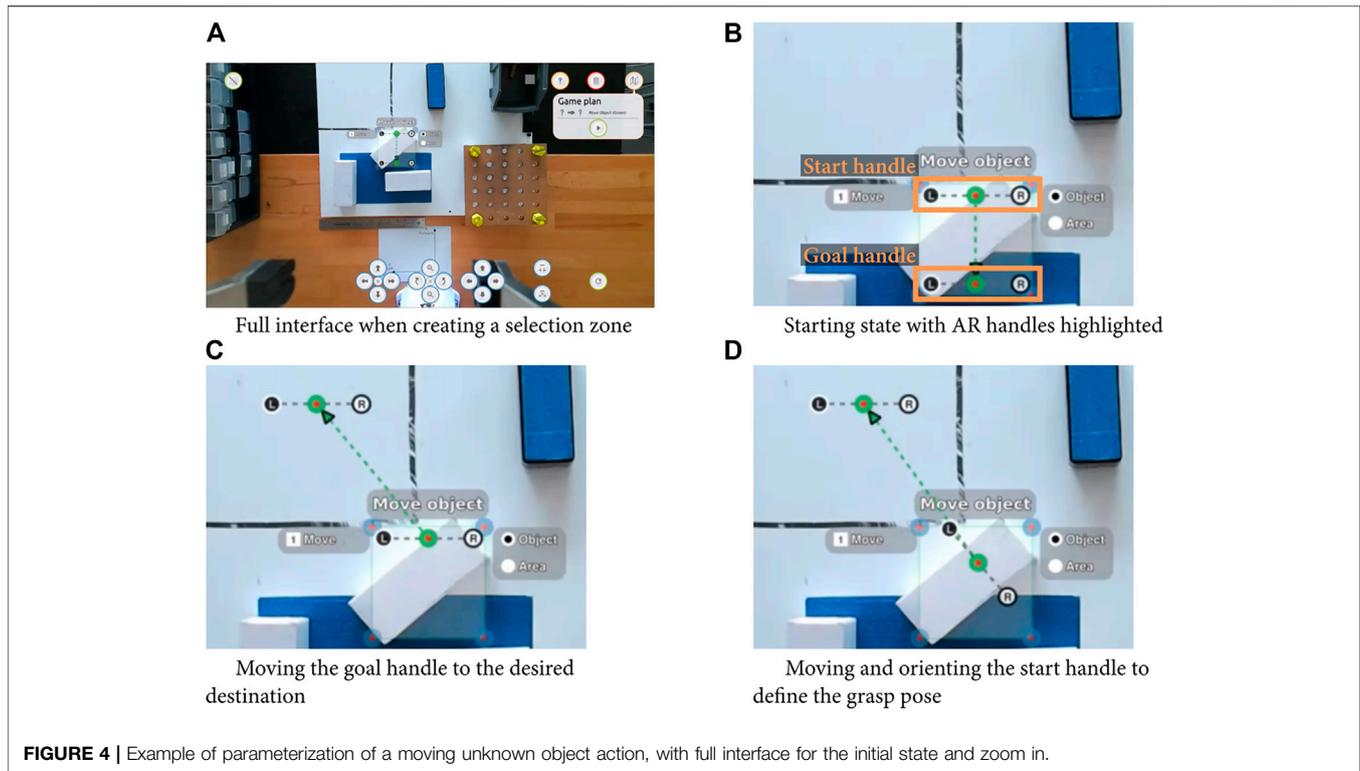
4.3.1 Direct Control

At the bottom of the screen there are a number of buttons for direct control: 12 buttons allow the user to move the camera by a discrete increment in each of the 6 potential directions (5 cm for the position buttons and $\pi/16$ radians for rotations), two buttons allow grasping and releasing, and a last button resets the robot to its homing position.

4.3.2 Authoring

To create task-level plans for the robots, users can annotate the augmented display to select actions applied to objects detected or parts of the environment. Users can click (or click-and-drag) on the screen to create selection areas to plan actions for the robot. Each selection area corresponds to one action or a set of actions on one type of object. Actions that can be parameterized (e.g., move actions) provide different types of handles that can be used to fully characterize the action. Users can create multiple selection areas to schedule different types of actions, and the resulting plan is shown in the Game Plan at the right of screen (see **Figure 2**). Users can use this game plan to confirm that the interface

¹Open-source code for our system implementation is available at <https://github.com/emmanuel-senft/authoring-ros/tree/study>.



interpreted the intentions correctly before sending the plan to the robot. During the execution, the user can monitor the robot progress in the task by watching the video feed and checking the progress in the plan. Video examples can be found at <https://osf.io/nd82j/>.

4.3.3 Interaction with undetected objects

To pick and place an object not detected by the system, users can manipulate a start pose and a goal pose handles to specify the motion (see **Figure 4**). These handles are composed of three connected points: the interaction point (grasping or releasing) as well as points representing the robot's fingers, and users can move the handle on the screen to change the interaction location, and rotate it to specify the end-effector orientation. This pixel value is then mapped into a 3D point in the camera frame using the Kinect's depth camera and converted in a point in space for robot. The orientation from the interface specifies the rotation on the vertical axis and consequently completely characterize a vertical tabletop grasp.

4.3.4 Generalization to Groups of Objects

When creating a selection area, the interface will select a default object to interact with based on the ones present in the area, but the type of object can be changed by clicking on radio buttons displaying the objects present in the area. Each object has a number of actions that can be executed on it (e.g., a screw can be tightened, loosened, or moved), and the user can select which action to apply and in what order by using numbered checkboxes. These actions will then be applied on each object of the selected type in the area (e.g., loosen and move all the screws in the area).

4.4 Backend

The interface exposed the following high-level actions: move (both known and unknown objects), loosen, tighten, wipe, pull, and push. These high-level actions selected by users are grounded in the real world by finding the 6D pose of the interaction points in the user plan (either using the depth map from the Kinect or the location of points in a list of known objects). Each action is then hierarchically decomposed in a set of lower level actions (e.g., pick-up, view, place) and primitives (e.g., move to position, move to contact, grasps). For example, a move known object action is decomposed first into a pick and a place actions, which are then decomposed into a multiple of primitives (move above grasping point, move to grasping point, grasp, move above grasping point, move above release point, move to release point, release, and finally move above releasing point). During execution, the robot will perform each of the parameterized primitive to complete the plan.

This method can be extended to new applications in three ways: 1) adding new objects to the image recognition and interface affordances, 2) by composing existing primitives to create new higher-level actions, and 3) if needed, by creating new primitives. The first two improvements could be made using graphical interface without having to code (e.g., using approaches similar to Steinmetz et al. (2018)), however the third one would require actual code modification. This is similar to the current state-of-the-art cobots teach pendants: they expose a number of primitives that users can use to create behaviors, but any requirement not covered by the primitives (such as additional sensor-based interaction) would need code development to add the capability. Nevertheless, we could envisage a mixed system

where creating new primitives and actions could be done locally, using learning from demonstration, and then exposed at a higher level to remote users using our interface.

5 EVALUATION

We conducted an evaluation to assess the impact of the design principles of our system. As mentioned in **Section 3**, our task-level authoring system is based around four principles: 1) interleaving observation and planning, 2) controlling the robot at the action level, 3) providing a unified, augmented reality interface, 4) graphical specification of actions. For the sake of the evaluation, the unified AR interface principle was not evaluated as too many different alternatives exist, however, we explored the three other axes. We conducted a 3×1 within-participants study to explore three types of interfaces embodying or not our design principles: our task-level authoring interface (TLA), a point-and-click interface (PC) inspired from Schmaus et al. (2019), and finally a Cartesian control interface (CC) as can be found traditionally on a cobot's teach pendant (e.g., PolyScope for the universal Robots²) or recent work in teleoperation (Marturi et al., 2016).

The CC condition does not use any of our design principles and serves as an alternative to kinesthetic teaching (Akgun et al., 2012) which cannot be applied due to the remote aspect and to direct control which would have required 6D input control on the user side. The PC condition only embodies the second design principle (controlling the robot at the action level). It corresponds to a simpler version of our interface, where the robot has similar manipulation capabilities (e.g., pick-up objects, loosen or tighten screws, pull drawers) but where actions can only be specified one at a time and where parameters have to be set numerically (e.g., using use sliders to specify parameters such as angles). The last condition TLA is the interface described in **Sections 3** and **4** and embodies all four of our principles.

The evaluation took place over Zoom,³ a video conference platform, and we use the built-in remote screen control as a way to allow participants to control the robot from remote locations. We did not assess the latency inherent of such system, but estimated it around one second.

5.1 Hypotheses

Our evaluation uses the metrics *S* for the task score, a performance measure; *A* for robot autonomy, measured by both total and individual periods of autonomy; *U* for usability, measured by the SUS scale (Brooke, 1996); *P* for user preference for the control method; and *W* for workload, measured by the NASA Task-Load Index (NASA TLX) (Hart and Staveland, 1988). Below, we describe our hypotheses and provide specific predictions for each measure. Subscripts denote study conditions (TLA, PC, and CC).

Our evaluation tested three hypotheses along the three evaluated design axes:

H1 Task score, autonomy, usability, and user preference will be higher, and workload will be lower with high-level control (PC, TLA) than low-level control (CC).

- Prediction P_{1a}: $S_{PC} > S_{CC}$, $A_{PC} > A_{CC}$, $U_{PC} > U_{CC}$, $P_{PC} > P_{CC}$, and $W_{PC} < W_{CC}$.

- Prediction P_{1b}: $S_{TLA} > S_{CC}$, $A_{TLA} > A_{CC}$, $U_{TLA} > U_{CC}$, $P_{TLA} > P_{CC}$, and $W_{TLA} < W_{CC}$.

H₂ Autonomy and user preference will be higher, and the workload will be lower when users are able to interleave observation and planning (TLA) than when they are not able to (PC).

- Prediction P_{2a}: $A_{TLA} > A_{PC}$.

- Prediction P_{2b}: $W_{TLA} < W_{PC}$.

- Prediction P_{2c}: $P_{TLA} > P_{PC}$.

H₃: Task score, usability, and user preference will be higher when users are able to parameterize actions graphically (TLA) than when they are not able to (PC).

- Prediction P_{3a}: $S_{TLA} > S_{PC}$.

- Prediction P_{3b}: $U_{TLA} > U_{PC}$.

- Prediction P_{3c}: $P_{TLA} > P_{PC}$.

H1 is based on the expectation that high-level action specification present in TLA and PC method automates away a large number of low-level actions that the user must specify in CC, which will save the user time and reduce the number of operations they must perform, thus their workload. H2 is grounded in the expectation that the task planning offered by our system will be used by participants to create longer periods of autonomy, which should reduce the workload, and make the participants prefer the method. Finally, H3 supposes that the graphical specification of actions will allow participants to specify action quicker (increasing their performance in the task), more easily (increasing the usability), and that participants will prefer this modality.

5.2 Method

5.2.1 Participants

We recruited 18 students enrolled in the Mechanical Engineering and Industrial and Systems Engineering departments at the university (3F/15M, age: $M = 19.6$, $SD = 1.54$). We selected our participants from this population as they represent people with some exposure but little expertise in robotics (familiarity with robots $M = 2.9$, $SD = 1.2$ on a five-point scale—none, a little, some, moderate, a lot—and familiarity with programming $M = 3$, $SD = 0.6$). The procedure was approved by the university's Institutional Review Board and participants were compensated at the rate of \$15/hour. The study was designed to last 80 min and included around 45 min of robot operation. Since it was completed remotely, participants stayed in their daily environment as shown in **Figure 5** top right, where a participant controlled the robot from his dorm bed.

5.2.2 Conditions

In all conditions, the layout of the interface was the same. It showed the camera feed, overlaid with arrows for direct control,

²<https://www.universal-robots.com>

³<https://zoom.us/>

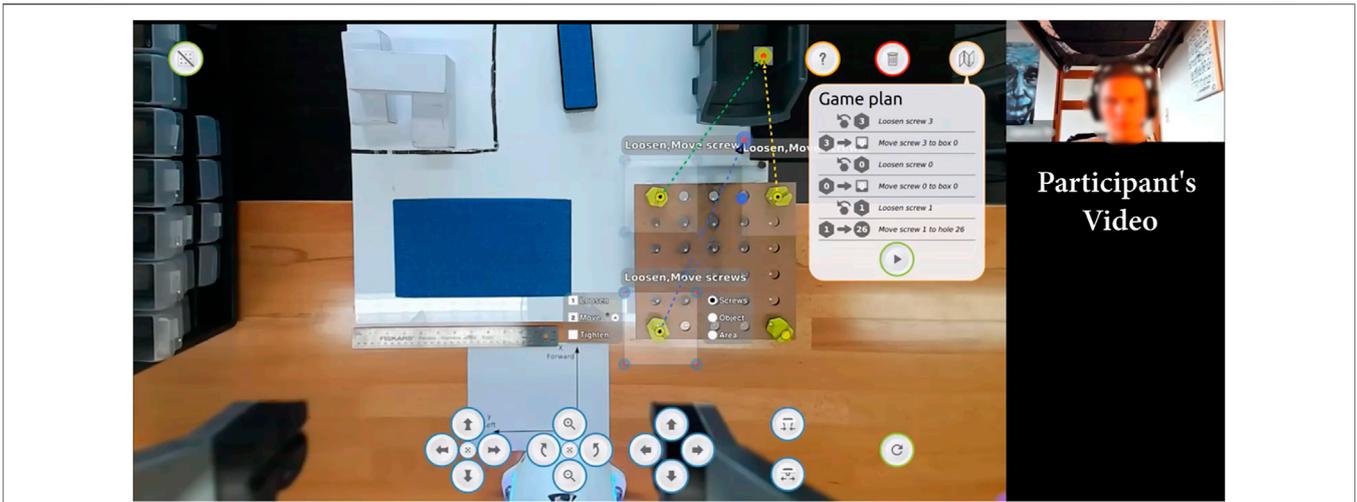


FIGURE 5 | Example of a participant using Drawing Board to control the robot from his dorm bed.



FIGURE 6 | Interfaces used in the study. Left shows the Cartesian Control interface: the user specifies numerically the end-effector position. Middle shows the Point-and-Click interface: the user selects direct actions on objects. Right shows the Task-Level Authoring interface allowing users to remotely create task plans for the robot.

buttons to grasp and release and the reset button. The difference was in the type of command sent to the robot as well as the modalities provided to the user. This study compared three conditions:

CC Cartesian control: the user uses six text boxes showing the current Cartesian position of the end-effector (x, y, z, rx, ry, rz) (see **Figure 6**-left). These text boxes can be edited with the desired command and sent to the robot either after modifying a single dimension or multiple ones.

PC Point-and-Click: the user is shown objects known by the system as markers overlaid on the video feed in an AR fashion (see **Figure 6**-center). The user can click on these markers or other parts of the view and is shown the different actions available on this object. Right clicking on an action allows to specify parameters, left clicking has the robot directly execute the action.

TLA Task-Level Authoring: interface presented in **Sections 3 and 4**, the user can annotate the video image to create actions associated to objects in the selection area and create task plans (see **Figure 6**-right).

To illustrate the difference between these conditions, we consider a move action on an unknown object (i.e., an object

the operator can see, but the robot does not identify). With the CC condition, participants had to enter the 6D pose of a grasping point. Often, this process would be iterative, the operation would first have the robot approach the object by specifying a higher point, then correct the position and angle, then move to the grasping point. Then the operator had to press the grasp button, move to a dropping point (by specifying the 6D pose or using the camera control buttons), then press release. With the PC conditions, participants could click on the grasping location on the screen, parameterize the action with a grasp angle, execute the pick-up action, reset the robot, click on the destination location on screen and select the place action. With TLA, participant could click on the screen to create a section area, keep the move unknown object action (the default one if no identified object was in the selection area), move the start and goal handles (as shown in **Figure 4**), and finally press execute.

5.2.3 Tasks

As shown in **Figure 6**, the workspace has a number of drawers on the left, three white boxes on the bottom, an eraser at the top and four screws on the right.

Participants had to complete a training task followed by four additional tasks:

Task 0: Training - move the angled white box to the top left area.

Task 1: Pick - and - place - move the additional two boxes (at different orientations) to the top left area.

Task 2: Repeated actions - loosen the four screws from the grid and move them to the top - right gray box.

Task 3: Exploration - locate a specific drawer on the left, pull it, inspect its content and push it back.

Task 4: Continuous action - wipe the blue area with the eraser.

These tasks were selected to represent different types of actions that a remote operator may need to complete. The first three pick-and-place actions free the area that need to be wiped, demonstrating workspace manipulation actions. The loosening and moving of the four screws represent repeated actions. The drawer inspection task combines two awareness acquisition actions: locating the drawer and inspecting its content, as well as two workspace manipulation actions: opening and closing the drawer. Finally, the wiping action represents a continuous action over an area, similar to cleaning a table or sanding a piece.

Of note, the three pick-and-place actions (task 0 and task 1) requires the human to specify manually the grasping and placing point as the robot does not detect the boxes by itself. And the exploration (task 3) requires the operator to gather information outside of the default field of view by moving the camera on the robot, read the labels on the drawers, locate the relevant drawer, open the drawer, look into the drawer, count the numbers of items, and finally close the drawer. To be able to complete this task autonomously, a robot would need to have optical character recognition capabilities and be able to detect and count any type of object present in the drawers. Furthermore, as shown in **Algorithm 1**, if an operator wanted to design in a single step a program solving this task, the resulting program would require logic functions such as loops, conditional on sensors, functions within conditions, and loop breaking conditions. All these functionalities could be supported by more complex visual programming languages (such as Blockly⁴) which requires more knowledge in programming. Such a program would also require more capabilities for the robot, more complex representation of the world (e.g., having a list of all drawers with positions to read the label from, and positions to inspect the content), and more complex programming languages. Instead, using a human-in-the-loop approach (through direct control or task-level authoring) allows to achieve the same outcome, but with much simpler robot capabilities and interfaces.

```

for Each drawers, d do
  goToReadingPosition(d);
  l = readLabel(d);
  if l == "Bolt M5" then
    pull(d);
    goToInspectPosition(d);
    countObjects(d, "Bolt M5");
    push(d);
    break;
  else
    continue;
  end
end

```

Algorithm 1 | Example of algorithm to solve the exploration task autonomously.

5.2.4 Procedure

Participants joined a zoom call from their home or other daily environment. The study started with informed consent and a demographic questionnaire. Then participants were asked to watch a video introducing the robot and the workspace, followed by a second video introducing the tasks participants would need to complete⁵. For each condition, participants first watched a 2-min video presenting the main modalities of the interface and demonstrating how to make a pick-and-place action. Then, participants had 15 min to complete as many of the tasks as possible. During the training, they could ask any questions to the experimenter, however in the four later tasks the experimenter was only able to answer the most simple questions (e.g., "the screws are tightened down, right?" but not "how to move this object?").

The interaction with the robot stopped when participants reached 15 min or when they completed all the tasks. After this interaction, participants filled out NASA Task-Load Index (NASA TLX) (Hart and Staveland, 1988) and System Usability Scale (SUS) (Brooke, 1996) questionnaires before moving on to the next condition. The order of the conditions was counter-balanced and the study concluded with a semi-structured interview and a debriefing where participants could ask questions to the experimenter. Despite our best efforts, some participants created actions that could trigger the robot's safety locks (often due to excessive force being applied). In such situations, the timer was paused, the robot was restarted and the participant continued from where they stopped.

⁴<https://developers.google.com/blockly>

⁵All the videos are available at <https://osf.io/nd82j/>.

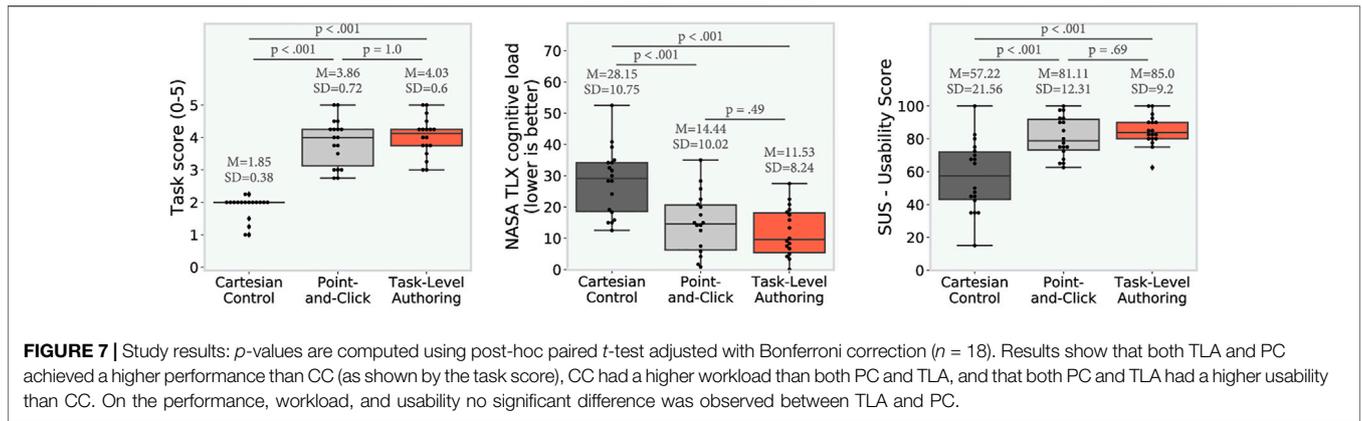


FIGURE 7 | Study results: *p*-values are computed using post-hoc paired *t*-test adjusted with Bonferroni correction (*n* = 18). Results show that both TLA and PC achieved a higher performance than CC (as shown by the task score), CC had a higher workload than both PC and TLA, and that both PC and TLA had a higher usability than CC. On the performance, workload, and usability no significant difference was observed between TLA and PC.

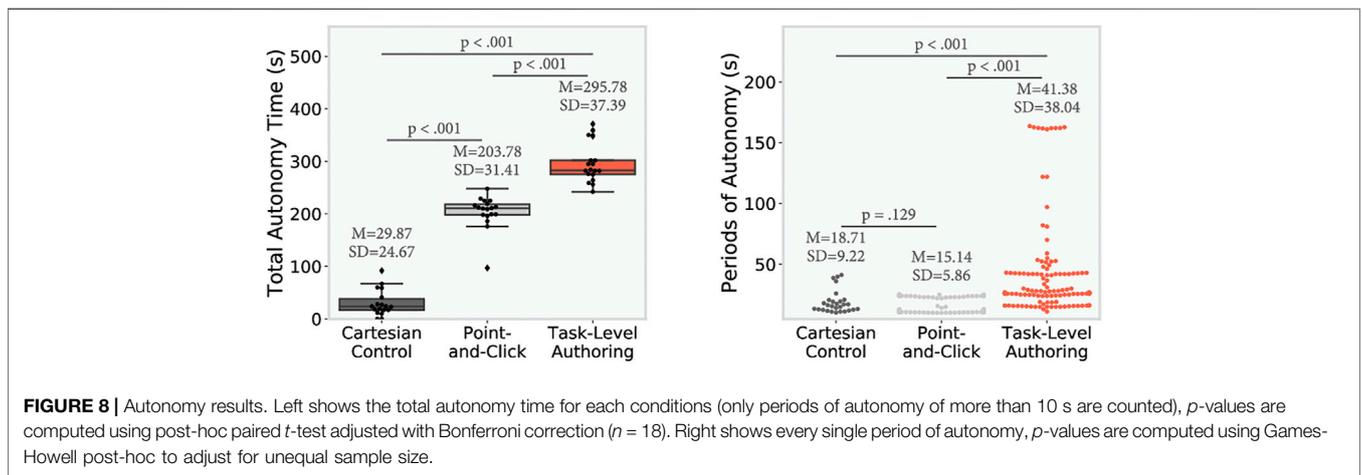


FIGURE 8 | Autonomy results. Left shows the total autonomy time for each conditions (only periods of autonomy of more than 10 s are counted), *p*-values are computed using post-hoc paired *t*-test adjusted with Bonferroni correction (*n* = 18). Right shows every single period of autonomy, *p*-values are computed using Games-Howell post-hoc to adjust for unequal sample size.

5.2.5 Measurements

We collected four types of quantitative data from the study. Task score, measured by how many tasks were fully or partially completed by the participants in the 15 min allocated per condition (with one point for the training and per task). Workload (i.e. how demanding it was to use the interface), measured by the NASA TLX. Usability (i.e. how intuitive the interface was), measured by SUS. Periods of autonomy, measured from period where the robot was moving continuously for more than 10 s (to be as inclusive as possible while not counting short periods that could barely be considered autonomous). We measured the periods of autonomy with a mixture of logs from the system and video coding of the interaction recordings using the Elan software (Nijmegen: Max Planck Institute for Psycholinguistics, The Language Archive, 2020).

In addition to the quantitative metrics, we collected qualitative impressions through the semi-structured interviews where we asked questions to the participants about their different experiences with the methods and which one they preferred.

6 RESULTS

Figure 7 and **Figure 8** present the quantitative results from the study. Results are first analyzed with a repeated measure ANOVA (corrected as needed using Greenhouse-Geisser), and then with post-hoc paired *t*-tests. A Bonferroni correction was directly applied to the *p*-values to protect against Type I error. For the periods of autonomy, as there was an unbalanced number of samples, we used ANOVA and Games-Howell post-hoc test.

6.1 Score

We observe significant impact of the condition on the score (sphericity was violated, Greenhouse-Geisser correction was used, $F(2, 34) = 115.53, p < 0.001$). Both the PC and the TLA interface achieved a score significantly higher than the CC (PC-CC: $t(17.0) = 13.0, p < 0.001$, TLA-CC: $t(17.0) = -23.0, p < 0.001$). However, we do not observe a significant difference of score between the PC and the TLA interfaces ($t(17.0) = -1.0, p = 1.0$).

Additionally, we did not observe an impact of the order, indicating that there no significant learning effect ($F(2, 34) = 0.698, p = 0.50$).

6.2 Periods of Autonomy

As shown in **Figure 8**, we observe a significant effect of the condition on the total autonomy time, $F(2, 34) = 297.45$, $p < 0.001$, and each condition was significantly different from the others, PC-CC: $t(17.0) = 18.0$, $p < 0.001$, TLA-CC: $t(17.0) = -25.0$, $p < 0.001$, TLA-PC: $t(17.0) = -7.0$, $p < 0.001$. With the CC offering the least amount of autonomy time, PC in the middle, and TLA offering the most of autonomy time. It can be observed that in addition to have a higher total autonomy time, the TLA condition also led to longer individual periods of autonomy than the other conditions (one way ANOVA: $F(2, 402) = 60.87$, $p < 0.001$, Games-Howell post-hoc test PC-CC: Mean Difference = -3.57 , $p = 0.129$, TLA-CC: Mean Difference = -22.66 , $p < 0.001$, TLA-PC: Mean Difference = -26.24 , $p < 0.001$).

6.3 Workload

We observe significant effect of the condition on workload as measured by the NASA TLX, $F(2, 34) = 46.29$, $p < 0.001$. Both the PC and the TLA interface imposed a workload significantly lower than the CC (PC-CC: $t(17.0) = -9.0$, $p < 0.001$, TLA-CC: $t(17.0) = 8.0$, $p < 0.001$). However, we do not observe a significant difference of workload between the PC and the TLA interfaces ($t(17.0) = 1.0$, $p = 0.49$).

6.4 Usability

We observe a significant effect of the condition on usability as measured by the SUS, $F(2, 34) = 23.18$, $p < 0.001$. Both the PC and the TLA interface were rated as having a high usability (SUS score around 80) significantly outperforming the Cartesian interface (PC-CC: $t(17.0) = 4.0$, $p < 0.001$, TLA-CC: $t(17.0) = -6.0$, $p < 0.001$). However, we do not observe a significant difference of usability between the PC and the TLA interfaces ($t(17.0) = -1.0$, $p = 0.69$).

6.5 Preference

When asked which methods they preferred, 14 participants replied they preferred the TLA method, three preferred the PC method and one the CC method. Using one-sample binomial test, we measure a significant preference for our TLA method (95% Adjusted Wald Confidence Interval is (54.24%, 91.54%), preference TLA >33% with $p < 0.001$).

6.6 Observations and Feedback From Participants

In addition of our quantitative metrics, we made a number of anecdotal observation during the study and the following semi-structured interview. First, the two main justifications for participant's preference of the TLA interface were the ability to queue actions and the visualization (the two design principles not supported by the PC interface):

"[TLA was] by far the best, because you could do so many tasks at once, and it was just really intuitive to figure out, okay, this is what it's gonna do"

"I like that little line thing which would show up on positions, so you could determine like initial position

and the final position [...] without having to remember numbers"

"Being able to angle the jaws, and have visual reference for that, was really useful"

Some participants used the periods of autonomy of the TLA interface to perform secondary actions, e.g., drinking water or even as one participant did, sending a message to a friend. Combined to our quantitative results showing that the authoring interface frees longer periods of time to the operators, these observations provide anecdotal evidence that interfaces similar to TLA could help operators perform secondary action. However, as our study did not assess such an hypothesis, future work should confirm it.

Some participants were slightly confused by the different modalities used and monitored the game plan to understand how their inputs were parsed:

"[TLA] also gave you that menu of like the order that you were going. I feel like that was really helpful"

Even though many participants qualified the PC interface of being very simple (almost too simple for some), participants still had to follow the progress of their series of action which can be complicated. For example, the screw task requires four repetitions of a loosen, a pick, a reset and a place action. Some participants in the study lost the track of which action was done, and for example forgot to loosen a screw, or did it twice. Some participants felt annoyed to have to re-specify each action every time:

"[With TLA] you could I guess perform multiple tasks at once, you didn't have to click on it every single times"

"[PC] it took more time, still because that you had to unscrew it, and then you had to pick it up, and then you had to move it and place it"

Being able to program the robots to sequence actions when having to repeat them over multiple objects allowed operators to keep track more easily of the progress in the task without having to keep in memory which actions were already executed.

Additionally, due to the lower granularity of control in CC and PC participants reported difficulties to know the distance between the robot's fingers and the table or objects or faced occlusion issues while operating the robot with the PC or CC interfaces. The task-level authoring offered by our system allowed participants to control the robot without facing these two obstacles.

7 DISCUSSION

7.1 Observations

Our results provide partial support for our hypotheses. H_1 is fully supported (both P_{1a} and P_{1b} are supported). The higher level interfaces performed better than Cartesian Control on all metrics supporting H_1 . H_2 is partially supported, TLA offered more autonomy than the other methods (supporting P_{2a}), and TLA

was preferred to the other methods and participants referred to the ability of queuing actions as a reason (supporting P_{2c}), however TLA did not reduce the workload compared to PC (failing to support P_{2b}). Finally, H_3 is also partially supporting. Participants did not achieve a higher score with TLA than with PC and did not rate the usability as higher for TLA (failing to support both P_{3a} and P_{3b}), however, participants did prefer TLA over PC and participants referred to the graphical action parameterization as a reason (supporting P_{3c}).

Overall, these results show that our design principles partially achieved their goals: the high-level control allowed participants to think at the task level and progress quicker in the tasks. TLA was preferred overall due to the opportunity to create flexible periods of autonomy and the graphical parameterization of actions. This flexible programming horizon allowed participants to specify long periods of autonomy when possible, but also directly select actions when the next step is unclear. Traditionally, robots with more autonomy will require a lower workload at runtime, as the operator does not need to provide inputs when the robot is autonomous. However, such autonomous robots might inflict a higher workload at design time and require more skills for the operator and more capabilities for the robot. By interleaving exploration, design of short plans, and execution, TLA aims to maintain this low workload both at runtime and design time. Compared to specifying a behavior a priori, allowing the operator to specify commands at runtime allows to solve similar problems, but with simpler robot capabilities (as the operators can perform some sensory analysis) and simpler interface (as the operators do not have to create programs handling every possible situation).

We observed a potential ceiling effect on the usability (a score of 85 on the SUS is defined as excellent usability (Brooke, 2013)), and possibly a floor effect on the workload (14 and 11 on the NASA TLX are very low scores). These effects may have two distinct origins. Either our study was not sufficiently challenging for our operators, or our action sequencing principle did allow participants to obtain capabilities closer to programming (through the scheduling of action, automatic generalization of a set of actions to a group of objects etc.), which may have increased the complexity of the interface, but our graphical specification principle balanced this added complexity to maintain a low workload and high usability. Due to time constraints and study complexity, we could not explore individually the impact of each axis, which prevents us to identify the root cause of this effect. Future work should investigate more precisely the situations in which these methods could differ in usability and workload.

Nevertheless, from our study we can confirm that our authoring interface allowed participants to specify longer plans for the robot and streamlined the execution of repeated and composite actions. These two additional benefits might allow operators to perform secondary tasks, and potentially facilitate extended use (as anecdotally supported by the observation that some participants lost track of their progress in the repeated action and did the same action twice, or forgot whether they unscrewed a bolt already). This additional gain comes at no cost in terms of workload and usability, which supports the conclusion that our design principles allowed the interface to be usable with limited training while incorporating additional programming capabilities. Future work should evaluate

whether such increase in autonomy could allow operators to perform secondary tasks in practice and how such programming capabilities could be used by operators.

7.2 Limitations

Our approach suffers from a number of limitations that we plan to address in future work. A key limitation is that the high-level interface requires the specific primitives and actions to be pre-determined and pre-programmed. Extending the set of operations to support a broader range of tasks may create challenges in helping the user understand the range of options. Allowing users to specify actions that are not in the interfaces “vocabulary” is challenging, as this requires detailed specification that often must consider low-level control issues such as compliance. This issue is common in authoring—for example, teach pendants are also intrinsically limited in the robot’s capabilities they expose and more complex uses often require coding. Additionally, our system relied on robust actions and we did not explore how to recover from failures when executing actions. We identify four ways such action failures could be handled. First, actions could be made more robust by integrating replanning strategies (e.g., planning a new grasp pose after a failed grasp). Second, high-level actions could take more parameters after a failure (e.g., specifying a full 6D grasp pose if the default one did not work). Third, the operator could provide additional runtime inputs to address small trajectory errors (Hagenow et al., 2021). And finally, the user could change the control mode for such infrequent event (e.g., temporarily using direct control instead of TLA).

The evaluation of our approach also has a number of limitations. For example, the study considered relatively simple tasks, used a mostly male population, and our population was not total novice, but had some experience with programming. Additionally, due to time constraints, we could not explore every single design axis individually. Future work should involve ablation studies, where the specific impact of design principles are evaluated, explore interactions in real environments, and use operators from the targeted population (family member controlling a robot in a home-assistant scenario or workers in a factory). Finally, future work should also explicitly explore the impact of latency when performing task-level authoring, especially compared to more direct control. We plan to address such limitations in future work.

7.3 Implications

Results from our evaluation lead to a number of implications. Centrally, the use of task-level authoring seems to be an interesting trade-off, allowing for sufficient programming to gain the advantages of asynchronous control (i.e., programming longer periods of autonomy for the robot and leading to longer and better quality idle time, offloading some tasks following to the robot), yet having the programming be simple enough that it can be used during the interaction with little training. The approach affords an interface design that combines exploration, specification, and monitoring in a single view. The specific interface provides other general lessons. First, our work expands on the ideas of using higher-level controls to enable effective teleoperation interfaces. While prior systems have shown point-and-click interfaces (Schmaus et al., 2019), ours expands the concept to accomplish longer autonomous behavior. Second, by connecting these higher-level controls in a paradigm where

exploration and manipulation are interleaved, we can create single-view interfaces that are usable in more complex scenarios. Third, our work extends prior see-through interfaces with camera control, allowing them to work in more environments. Fourth, our work shows the potential of asynchronous interfaces by improving the amount and duration of the offered periods of autonomy. By allowing the user to quickly specify longer plans, they gain opportunities for idle time, potentially freeing them to perform other tasks during execution. Finally, by demonstrating effective telemanipulation only using consumer interfaces shows that remote robot operation is possible for novice users—even at distances of many time zones.

8 CONCLUSION

In this paper, we explored the design of interfaces for remote control of a robotic arm by novice users. Our design considers the key goals of teleoperation interfaces: allowing remote novice operators to analyze the robot's environment and specify robot behavior appropriate to the situation. To address these challenges for scenarios with novice users and standard input devices we adopted a task-level authoring approach. The approach allowed for the design of an interface that interleaves exploration and planning, allowing us to utilize both direct control (more intuitive interface and benefiting from the human knowledge more frequently) and asynchronous control (robustness to communications issues and increased idle time for the operator). Our interface uses graphical overlays on a video feed of the environment to provide for simple exploration, specification of operations, and sequencing of commands into short programs. We evaluated a prototype system in an 18-participant study which showed that our interface allowed users with some familiarity with programming to 1) operate the robot remotely to gain awareness about the environment, 2) perform manipulation of the workspace, and 3) use the scheduling of actions to free long periods of idle times that might be used to perform secondary tasks. Furthermore, our interface was largely preferred compared to two other simpler interfaces.

REFERENCES

- Akan, B., Ameri, A., Cürüklü, B., and Asplund, L. (2011). Intuitive Industrial Robot Programming through Incremental Multimodal Language and Augmented Reality. In 2011 IEEE International Conference on Robotics and Automation (IEEE), 3934–3939. doi:10.1109/icra.2011.5979887
- Akgun, B., Cakmak, M., Yoo, J. W., and Thomaz, A. L. (2012). Trajectories and Keyframes for Kinesthetic Teaching: A Human-Robot Interaction Perspective. In Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction. 391–398.
- Azuma, R. T. (1997). A Survey of Augmented Reality. *Presence: Teleoperators & Virtual Environments* 6, 355–385. doi:10.1162/pres.1997.6.4.355
- Bajcsy, R. (1988). Active Perception. *Proc. IEEE* 76, 966–1005. doi:10.1109/5.5968
- Beer, J. M., Fisk, A. D., and Rogers, W. A. (2014). Toward a Framework for Levels of Robot Autonomy in Human-Robot Interaction. *J. human-robot interaction* 3, 74. doi:10.5898/jhri.3.2.beer
- Billard, A., Calinon, S., Dillmann, R., and Schaal, S. (2008). *Survey: Robot Programming by Demonstration*. Berlin: Springer Handbook of Robotics. doi:10.1007/978-3-540-30301-5_60

Our work adds a new tool to the existing library of teleoperation approaches and demonstrates that task-level authoring is a powerful method to allow non-experts to remotely create short periods of autonomy for robots while allowing them to explore the robot's environment.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by the University of Wisconsin-Madison Institutional Review Board. The patients/participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

ES, BM, and MG contributed to conception and design of the study. ES, KW, and MH designed and implemented the prototype. ES ran the study and performed the statistical analysis. ES wrote the first draft of the manuscript. MH, KW, MG, and BM wrote sections of the manuscript. BM, MG, RR, and MZ obtained funding for the project. All authors contributed to manuscript revision, read, and approved the submitted version.

FUNDING

This work was supported by a NASA University Leadership Initiative (ULI) grant awarded to the UW-Madison and The Boeing Company (Cooperative Agreement # 80NSSC19M0124).

- Bohren, J., Papazov, C., Burschka, D., Krieger, K., Parusel, S., Haddadin, S., et al. (2013). A Pilot Study in Vision-Based Augmented Telemanipulation for Remote Assembly over High-Latency Networks. In 2013 IEEE International Conference on Robotics and Automation. IEEE, 3631–3638. doi:10.1109/icra.2013.6631087
- Brooke, J. (2013). Sus: a Retrospective. *J. usability Stud.* 8, 29–40.
- Brooke, J. (1996). Sus: a “Quick and Dirty” usability. *Usability Eval. industry*, 189.
- Cao, Y., Wang, T., Qian, X., Rao, P. S., Wadhawan, M., Huo, K., et al. (2019a). Ghostar: A Time-Space Editor for Embodied Authoring of Human-Robot Collaborative Task with Augmented Reality. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology. 521–534. doi:10.1145/3332165.3347902
- Cao, Y., Xu, Z., Li, F., Zhong, W., Huo, K., and Ramani, K. (2019b). V. Ra: An In-Situ Visual Authoring System for Robot-Iot Task Planning with Augmented Reality. In Proceedings of the 2019 on Designing Interactive Systems Conference. 1059–1070.
- Chen, H., Kakiuchi, Y., Saito, M., Okada, K., and Inaba, M. (2011). View-based Multi-Touch Gesture Interface for Furniture Manipulation Robots. In Advanced Robotics and its Social Impacts. IEEE, 39–42. doi:10.1109/arsos.2011.6301979
- Datta, C., Jayawardena, C., Kuo, I. H., and MacDonald, B. A. (2012). Robostudio: A Visual Programming Environment for Rapid Authoring and Customization of Complex

- Services on a Personal Service Robot. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, 2352–2357. doi:10.1109/iros.2012.6386105
- Funda, J., Lindsay, T. S., and Paul, R. P. (1992). Teleprogramming: Toward Delay-Invariant Remote Manipulation. *Presence: Teleoperators & Virtual Environments* 1, 29–44. doi:10.1162/pres.1992.1.1.29
- Gao, Y., and Huang, C.-M. (2019). Pati: a Projection-Based Augmented Table-Top Interface for Robot Programming. In Proceedings of the 24th international conference on intelligent user interfaces. 345–355.
- Glas, D. F., Kanda, T., and Ishiguro, H. (2016). Human-robot Interaction Design Using Interaction Composer Eight Years of Lessons Learned. In 2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 303–310. doi:10.1109/hri.2016.7451766
- Guerin, K. R., Lea, C., Paxton, C., and Hager, G. D. (2015). A Framework for End-User Instruction of a Robot Assistant for Manufacturing. In 2015 IEEE international conference on robotics and automation (ICRA). IEEE, 6167–6174. doi:10.1109/icra.2015.7140065
- Hagenow, M., Senft, E., Radwin, R., Gleicher, M., Mutlu, B., and Zinn, M. (2021). Corrective Shared Autonomy for Addressing Task Variability. *IEEE Robot. Autom. Lett.* 6, 3720–3727. doi:10.1109/lra.2021.3064500
- Hart, S., Dinh, P., and Hambuchen, K. (2014). Affordance Templates for Shared Robot Control.
- Hart, S. G., and Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Adv. Psychol. (Elsevier)* 52, 139–183. doi:10.1016/s0166-4115(08)62386-9
- Hashimoto, S., Ishida, A., Inami, M., and Igarashi, T. (2011). Touchme: An Augmented Reality Based Remote Robot Manipulation. In The 21st International Conference on Artificial Reality and Telexistence, Proceedings of ICAT2011. vol 2.
- Huang, J., and Cakmak, M. (2017). Code3: A System for End-To-End Programming of mobile Manipulator Robots for Novices and Experts. In 2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 453–462.
- Johnson, M., Shrewsbury, B., Bertrand, S., Wu, T., Duran, D., Floyd, M., et al. (2015). Team IHMC's Lessons Learned from the DARPA Robotics Challenge Trials. *J. Field Robotics* 32, 192–208. doi:10.1002/rob.21571
- Kent, D., Saldanha, C., and Chernova, S. (2020). Leveraging Depth Data in Remote Robot Teleoperation Interfaces for General Object Manipulation. *Int. J. Robotics Res.* 39, 39–53. doi:10.1177/0278364919888565
- Marescaux, J., Leroy, J., Gagner, M., Rubino, F., Mutter, D., Vix, M., et al. (2001). Transatlantic Robot-Assisted Telesurgery. *Nature* 413, 379–380. doi:10.1038/35096636
- Marturi, N., Rastegarpanah, A., Takahashi, C., Adjigble, M., Stolkin, R., Zurek, S., et al. (2016). Towards Advanced Robotic Manipulation for Nuclear Decommissioning: A Pilot Study on Tele-Operation and Autonomy. In 2016 International Conference on Robotics and Automation for Humanitarian Applications (RAHA). IEEE, 1–8. doi:10.1109/raha.2016.7931866
- Maxwell, S., Cooper, B., Hartman, F., Leger, C., Wright, J., and Yen, J. (2005). The Best of Both Worlds: Integrating Textual and Visual Command Interfaces for mars Rover Operations. In 2005 IEEE International Conference on Systems, Man and Cybernetics (IEEE). 2, 1384–1388.
- Nakaoka, S., Morisawa, M., Kaneko, K., Kajita, S., and Kanehiro, F. (2014). Development of an Indirect-type Teleoperation Interface for Biped Humanoid Robots. In 2014 IEEE/SICE International Symposium on System Integration. IEEE, 590–596.
- Niemeyer, G., Preusche, C., Stramigioli, S., and Lee, D. (2016). Telerobotics. In Springer handbook of robotics. Springer, 1085–1108. doi:10.1007/978-3-319-32552-1_43
- Nijmegen: Max Planck Institute for Psycholinguistics, The Language Archive (2020). Elan (version 5.9) [computer software].
- Norris, J. S., Powell, M. W., Fox, J. M., Rabe, K. J., and Shu, I.-H. (2005). Science Operations Interfaces for mars Surface Exploration. In 2005 IEEE International Conference on Systems, Man and Cybernetics, 2. IEEE, 1365–1371.
- Paxton, C., Hundt, A., Jonathan, F., Guerin, K., and Hager, G. D. (2017). Costar: Instructing Collaborative Robots with Behavior Trees and Vision. In 2017 IEEE international conference on robotics and automation (ICRA). IEEE, 564–571. doi:10.1109/icra.2017.7989070
- Peng, H., Briggs, J., Wang, C.-Y., Guo, K., Kider, J., Mueller, S., et al. (2018). Roma: Interactive Fabrication with Augmented Reality and a Robotic 3d Printer. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. 1–12.
- Perzly, A., Somani, N., Profanter, S., Kessler, I., Rickert, M., and Knoll, A. (2016). Intuitive Instruction of Industrial Robots: Semantic Process Descriptions for Small Lot Production. In 2016 IEEE/RSJ international conference on intelligent robots and systems (iros). IEEE, 2293–2300. doi:10.1109/iros.2016.7759358
- Porfirio, D., Fisher, E., Sauppé, A., Albarghouthi, A., and Mutlu, B. (2019). Bodystorming Human-Robot Interactions. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology. 479–491. doi:10.1145/3332165.3347957
- Preusche, C., Reintsema, D., Landzettel, K., and Hirzinger, G. (2006). Robotics Component Verification on ISS ROKVISS-Preliminary Results for Telepresence. In 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, 4595–4601. doi:10.1109/iros.2006.282165
- Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., et al. (2009). Ros: an Open-Source Robot Operating System. In ICRA workshop on open source software, 3. Japan: Kobe, 5.
- Rakita, D., Mutlu, B., and Gleicher, M. (2018). An Autonomous Dynamic Camera Method for Effective Remote Teleoperation. In Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. 325–333. doi:10.1145/3171221.3171279
- Rosenberg, L. B. (1993). Virtual Fixtures: Perceptual Tools for Telerobotic Manipulation. In Proceedings of IEEE virtual reality annual international symposium. IEEE, 76–82. doi:10.1109/vrais.1993.380795
- Schmaus, P., Leidner, D., Krüger, T., Bayer, R., Pleintinger, B., Schiele, A., et al. (2019). Knowledge Driven Orbit-To-Ground Teleoperation of a Robot Coworker. *IEEE Robotics Automation Lett.* 5, 143–150.
- Sefidgar, Y. S., Agarwal, P., and Cakmak, M. (2017). Situated Tangible Robot Programming. In 2017 12th ACM/IEEE International Conference on Human-Robot Interaction HRI. IEEE, 473–482. doi:10.1145/2909824.3020240
- Steinmetz, F., Wollschläger, A., and Weitschat, R. (2018). RAZER-A HRI for Visual Task-Level Programming and Intuitive Skill Parameterization. *IEEE Robot. Autom. Lett.* 3, 1362–1369. doi:10.1109/lra.2018.2798300
- Walker, M. E., Hedayati, H., and Szafir, D. (2019a). Robot Teleoperation with Augmented Reality Virtual Surrogates. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, 202–210. doi:10.1109/hri.2019.8673306
- Walker, N., Peng, Y.-T., and Cakmak, M. (2019b). *Robot World Cup*. Springer, 337–350. doi:10.1007/978-3-030-35699-6_26
- Neural Semantic Parsing with Anonymization for Command Understanding in General-Purpose Service Robots
- Weintrop, D., Afzal, A., Salac, J., Francis, P., Li, B., Shepherd, D. C., et al. (2018). Evaluating Coblox: A Comparative Study of Robotics Programming Environments for Adult Novices. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. 1–12.
- Wu, Y., Kirillov, A., Massa, F., Lo, W.-Y., and Girshick, R. (2019). Detectron2. available at: <https://github.com/facebookresearch/detectron2>.
- Yamauchi, B. M. (2004). Packbot: a Versatile Platform for Military Robotics. In Proc. SPIE 5422, Unmanned Ground Vehicle Technology VI, September 2, 2004, 228–237. doi:10.1117/12.538328
- Yanco, H. A., Norton, A., Ober, W., Shane, D., Skinner, A., and Vice, J. (2015). Analysis of Human-Robot Interaction at the Darpa Robotics Challenge Trials. *J. Field Robotics* 32, 420–444. doi:10.1002/rob.21568
- Yanco, H., Drury, J., and Scholtz, J. (2004). Beyond Usability Evaluation: Analysis of Human-Robot Interaction at a Major Robotics Competition. *Human-comp. Interaction* 19, 117–149. doi:10.1207/s15327051hci1901&2_6

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Senft, Hagenow, Welsh, Radwin, Zinn, Gleicher and Mutlu. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.